(2)

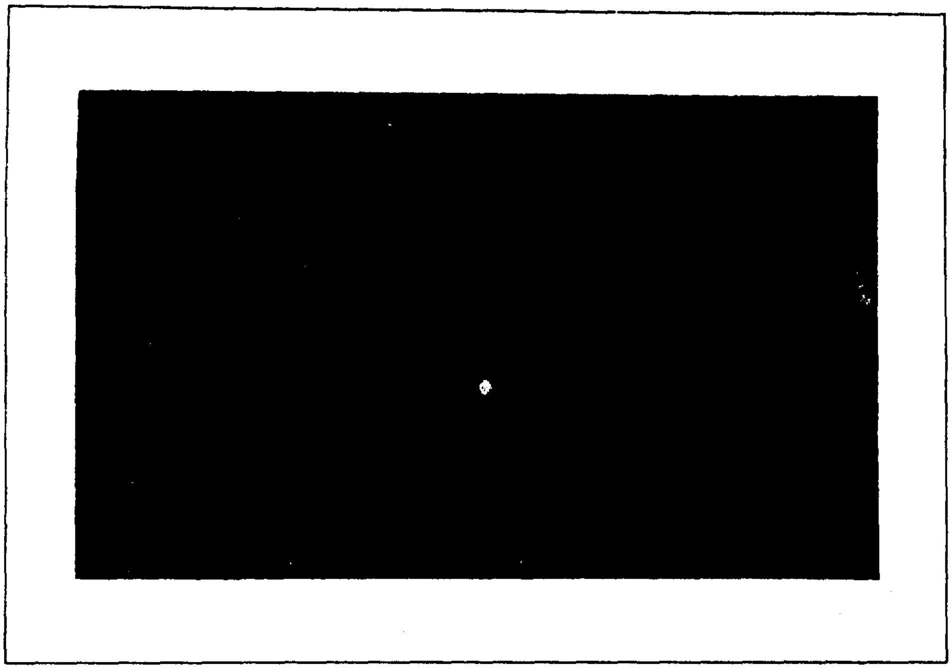

The Artificial Intelligence and Psychology Project

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

(2)

Qualitative Reasoning: How We Think Our Way Through the Day

Technical Report AIP - 130

Herbert A. Simon

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

March 1, 1990

This research was supported in part by the Defense Advanced Research Projects Agency under Contract F-33615-84-K-1520, and in part by the Personnel and Training Programs, Computer Sciences Division, Office of Naval Research, under contract number N00014-86-K-0678.

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; Distribution unlimited
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) 130	5. MONITORING ORGANIZATION REPORT NUMBER(S)
6a. NAME OF PERFORMING ORGANIZATION Carnegie Mellon University 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research (Code 1133)
6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, PA 15213	7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, VA 22217-5000
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Same as Monitoring Organization 8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678
8c. ADDRESS (City, State, and ZIP Code)	10. SOURCE OF FUNDING NUMBERS p40005ub201/7-4-86 PROGRAM ELEMENT NO. N/A; PROJECT NO. N/A; TASK NO. N/A; WORK UNIT ACCESSION NO. N/A

11. TITLE (Include Security Classification)

QUALITATIVE REASONING: HOW WE THINK OUR WAY THROUGH THE DAY

12. PERSONAL AUTHOR(S) Herbert A. Simon

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Technical	FROM 86Sept15 TO 91Sept14	90/3/1	13

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	qualitative reasoning; inference from diagrams; mental imagery; expert systems; heuristic reasoning; monotonic transformations

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

SEE REVERSE SIDE

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Meyrowitz	22b. TELEPHONE (Include Area Code) (202) 696-4302 22c. OFFICE SYMBOL N00014

ABSTRACT

We have come full circle back to the topic of visualization and the role of the mind's eye and of external visual displays in human thinking. I have tried to survey some of the main tools and processes that seem to be implicated in everyday reasoning -- the kind that carries us through the day, dealing with problems as they arise.

The thinking I have described does not look at all like formal logic, and only a little like mathematics. It makes use of a great multitude of inference rules, which are not tautological rules of logic but incorporate much real-world knowledge. It appears to be remarkably unconcerned with questions of sufficiency and necessity. When it deals with quantities, as it often must, it usually handles primarily their ordinal rather than their cardinal properties. For most people, at least, it makes great use of diagrammatic representations, or mental diagrams in the mind's eye, which provide it with pwoerful inference processes. To compensate for its severe limitations in handling simultaneous relations, it proceeds by successive approximations, and halts when it has satisficed.

By the standards of formal logic, it is a jerry-built structure. But it gets us through the day.

Qualitative Reasoning: How we Think our Way Through the Day

Herbert A. Simon
Carnegie-Mellon University

The term "reasoning" has been co-opted by logicians, experimental psychologists, and perhaps others to refer almost exclusively to deductive and inductive inference. In my remarks here, I intend to use the term much more broadly -- essentially as coextensive with thinking. It may be true, but it is certainly not obvious, that all thinking is reasoning in the logician's strict sense. For example, we ordinarily regard the use of analogy and metaphor as thinking, but these processes do not fit very comfortably the logicians' paradigms for reasoning.

Reasoning by "Seeing"

Consider an example even more remote from logic. I ask you to visualize a rectangle in the vertical plane that is twice as wide as it is high. I now ask you to drop a vertical line from the midpoint of the upper edge to the midpoint of the lower edge. Next I ask you to draw a diagonal from the northwest corner of the rectangle to the southeast corner. Does this diagonal intersect the vertical line you drew previously?

Almost everyone will agree that it does -- in fact, will be very sure that it does. But how do you know? Did you prove it, using the axioms of Euclidean geometry or an algebraic calculation based on analytic geometry? Probably not. You most likely just *saw* it.

I am going to define reasoning broadly enough to encompass inferences like the one we have just made. We represent information in various ways -- in this case, in the form of a mental picture.[1] We have available certain processes, or operators -- some of them conscious, some not -- for operating on this information and drawing conclusions from it. In the case at hand, the subconscious operator we applied is one we call "seeing." It is available for extracting information both from external visible displays and from pictures in the "mind's eye."

Inferring something because you can see it in your mind's eye, or even on a piece of paper, is not always a reliable process. In either case, there are severe limits on the resolution of the picture and the accuracy of the information-extracting processes. Consider the following example. We draw a square with one inch sides. Around the northwest corner of the square as

[1]By "mental picture" I mean whatever representation of the rectangle in your head permitted you to infer that the diagonal and vertical lines intersected.

center, we draw a circle with radius equal to the side of the square. Around the southeast corner we draw a second circle with radius the length of half the side.

Do the two circles intersect? Can you see the intersection, or the lack of one? How certain are you of your answer? A fairly simple calculation, which I leave to you, will convince you that the two circles do indeed intersect. But if the calculation convinces you, remember that it is a calculation and not a "seeing." What are the limits of what you or I can see in a simple geometric figure?

Formal Reasoning

A more or less justified claim that is often made for reasoning that employs the strict rules of formal logic is that it is perfectly reliable. A modern system of formal logic provides you with a small set of axioms and an even smaller set of processes (inference rules) for deriving new expressions from the axioms. In *Principia Mathematica*, for example, there are about a half dozen axioms and only two inference rules, *modus ponens* and substitution. Every derivation consists of a sequence of easily checked steps, for the terminal expression in each step must be derivable, applying one of the two rules of inference, from axioms or theorems previously derived.

In spite of the reliability that should be obtainable in this way, logic suffers from grave problems that begin at the very foundations. First, there are the Gödel theorems, which guarantee, for an even moderately rich logic, that if the logic is consistent (and who would want an inconsistent logic?) it is incomplete: there are true statements expressible in the language of the logic that are not provable. Moreover, as such distinguished logicians as Frege, Russell, and Quine have at one time or another discovered, a contradiction may actually be lurking in even the most carefully crafted logic.

But logic suffers from even graver difficulties than the possible presence of contradictions and the certain presence of incompleteness. Just because it restricts itself to a small number of inference rules, it proceeds by tiny steps, hence takes almost forever to get anywhere. Almost no serious mathematics is done using the strict rules of the game of mathematical logic.

Working mathematicians use seven-league boots to take giant steps from already proved expressions to new consequences. They do so with a certain measure of confidence that anyone skilled in the art, given time and patience, can fill in the numerous missing intermediate steps. Often they are wrong (witness the recent failed claim of a proof for Fermat's Last Theorem), and have to patch up their proofs by restricting the claims of their theorems or in other ways. If they are good mathematicians, and not too adventuresome, their proofs are right more often than not, but there is nothing like absolute certainty in them.

Heuristic Reasoning

It would appear, then, that mathematicians reason just like other people. They are not a separate subspecies of the human race. They acquire, in the course of their training and experience, a rich set of inference operators that allow them, with a certain degree of assurance, to proceed from a description of some state of affairs to inferences about it that are implied but not explicitly stated in the description. Their inference processes provide a certain level of reliability in their reasoning, but not certainty. In return for taking the risk of being mistaken, mathematiciaans are rewarded by being able to reason far more powerfully and rapidly than if they stuck to the strict rules of the game of logic. They can do serious mathematics, and are not restricted to trivial mathematics.

The prudent among them, and among those of us who are not mathematicians, often divide the reasoning process into two parts, a discovery part and a verification part. First, we rely on quite powerful but somewhat unreliable inference operators, so-called *heuristic processes*, to discover new conclusions. Then, before we announce our new truths or act on them, we subject them to careful scrutiny, using more microscopic and reliable inference rules to test the doubtful steps. In the common practice of reasoning, discovery and verification are intermingled, for we may alternate between them during the whole course of our thinking.

If absolute reliability is not essential, we can employ even more powerful, but less reliable, heuristic rules, and be quite relaxed about verification. We may be obliged to do so when we have to act in real time -- need we verify in a strictly logical way that the automobile is about to run over us? Or we may prefer to use unreliable but powerful procedures if mistakes can be repaired, are not irrevocable. Then we leave verification to Nature, who will ultimately tell us if we were in error.

The conclusion we have reached is that human reasoning cannot be identified with the processes prescribed in books on formal logic. It is much less preoccupied with absolute reliability than logic is, and much concerned with the availability of inference operators that are powerful, if sometimes fallible.

Thinking is not Tautological

One characteristic of logic I have not focused upon is that formal logic is generally supposed to be tautological. That is to say, the conclusions of an argument in formal logic should not place any restriction on the range of possible worlds that are not imposed by the axioms. If the premises admit unicorns, then the conclusions must not exclude them.

It is because of this tautological or analytic character of the inference rules of a proper logic that we have confidence in the conclusion of a chain of reasonings without feeling the need of checking it by empirical observation or experiment. If we do make empirical tests of it and the tests fail, we do not blame the inference rules but the axioms.

"All men are mortal. Socrates is a man. Therefore, Socrates is mortal." We do not question the structure of the syllogism. If we find that Socrates is still alive today, we conclude either that his death will occur at some future date or that the major or minor premises of the syllogism are in error. The rule of *modus ponens* is irreproachably analytic, and cannot be invalidated by empirical evidence.

But do we follow this principle in our everyday reasoning? A weight is hanging from one end of a rope thrown over a pulley, and a second weight from the other end of the rope. Both weights are stationary, and the pulley looks quite new and not rusted. We notice that the weight attached to the left end of the rope is ten kilograms, and conclude that the weight attached to the right end is also ten kilograms.

Perhaps we reasoned to this conclusion using only analytic inference rules. Perhaps all of our knowledge of the laws of physics was stored in axioms. It would be an interesting exercise to construct a proof along these lines, using only *modus ponens* and substitution as rules of inference. But I cannot wait until you have finished it. What is the alternative?

The alternative is to include laws of physics, not just among the axioms, but among the inference rules as well. We can write the rules in the form of *productions*, C--->A, which read "Whenever the conditions *C* are satisfied, execute the actions *A*." In the case before us, the production might (roughly) take the form: "If you see a weight hanging from the rope on one side of a pulley wheel, and another weight on the other side, and if the pulley is frictionless, and if the weights are motionless, and if the weight on one side has a known value ---> assign the same value to the weight on the other side."

Of course you pay a price for this law of inference. It will lead you to correct conclusions only in a world governed by Newton's Laws of Motion (and only in a world where you can detect reliably whether a pulley is frictionless). If the conclusion, subjected to other tests, turns out to be false, it may be your inference rule that is at fault, and not your other assumptions.

Note that in proceeding in this way, we are doing more than just assuming Newton's Laws of Motion, or even embodying them in inference rules. We are giving ourselves powerful means of making inferences that would otherwise be very difficult to make. Suppose, for example, that the pulley problem is presented to us as a diagram in a physics book, or that it is presented in words, but we form our own mental diagram. Then if we happen to have acquired a production corresponding to Newton's Third Law, we will be able to reason about the pulley.

What does it mean "to have acquired a production"? It means that when a pulley presents itself satisfying the conditions of the production, we will notice it and that the conditions are satisfied, the corresponding actions will be evoked from our memory, and we will take the action -- in this case, reach a conclusion about the unknown weight.

Inference Chains

Our experiences in geometry or logic or in mathematics generally may lead us to think of reasoning as a process of *constructing* long chains of steps that lead from premises to conclusion. In everyday life, reasoning is seldom quite like that. In any specific situation within a given task, only a few productions are likely to be evoked at any moment.

Human attention at any given time focuses on only a very restricted set of stimulus features, visible or audible in the external environment, or recovered from memory and held temporarily in short-term memory. It is this limited set of stimuli, noticed and held in attention, that must supply the cues that evoke productions -- the conditions that must be satisfied. The evokable productions may be further restricted to those that are relevant to the goal at hand, for one of the conditions a production may depend on is a goal symbol held in short-term memory. In the absence of that symbol, the production's conditions will not be satisfied and the production will not be evoked.

For example, consider the productions used by someone with the requisite skills to solve a simple equation in algebra. Memory holds the goal of solving the equation. If a constant term is noticed to the left of the equals sign, a production is evoked that substracts the term from both sides of the equation. If a term in the variable is noticed to the right of the equals sign, a second production is evoked that substracts that term from both sides of the equation. If the variable term on the left of the equals sign has a coefficient other than unity, a third production divides all terms in the equation by that coefficient.

Other than collecting terms, that's all there is to solving such an equation. What is required is that the three productions described above be available, that they be evoked by the cues on the paper and in memory, and that they not be interfered with by a lot of other irrelevant productions evoked at the same time. The latter obstruction is prevented by the narrow span of attention and the presence of appropriate goal symbols among the conditions of the productions.

It should be noted especially that the action of each production creates a new situation that makes the action of the next production appropriate and at the same time evokes it. The chaining is accomplished by the continually changing problem situation that evokes the right chain of productions. The productions do not have to be strung together in memory. The situation evokes them at the proper time.

The scheme I have just described provides the backbone for most of those reasoning systems that we call "expert systems." A medical diagnostic system, for example, responds to symptoms by hypothesizing disease entities. The presence of such hypotheses evokes other productions that propose additional tests to discriminate among alternatives. After a diagnosis has been accepted, yet other productions will supply a Latin name for the disease entity, will make a prognosis, will propose treatment, and so on. The result may be "logical" enough, but there is little about the process that suggests the long chains of steps that characterize formal logical inference. At each successive moment, the system simply does "what the situation calls for," or even more accurately, "what it is reminded of by the situation."

Each production can, of course, be formally regarded as an inference rule. If we tried to state just what the inference is, we would have to say something like, "The situation satisfying conditions C implies that it would be a good idea to take the actions A." In the pulley example, it being known (and noticed) that the force on the left end of the rope over the pulley is F, it is a good idea to assign the force F to the right end of the rope as well. The implication is a very conditional one, requiring both that our goal is to find the forces in the pulley system and that the system is governed by Newton's Laws. If either condition is not satisfied, the production is inappropriate, and we would hope that it would not be evoked.

A crucial difference between the productions and usual rules of inference is that the former are mandatory (they are to be executed whenever their conditions are satisfied), while the latter are permissive (they may be executed at any time, but it is nowhere specified under what circumstances they *should* be executed).

Necessity and Sufficiency

In formal logic and mathematics we pay a great deal of attention to whether our premises are necessary and/or sufficient for our conclusions. Insufficiency is, of course, the mortal sin, since it means that our conclusions need not hold under all conditions. Lack of necessity is only a venal sin, an unesthetic redundancy that makes us fall short of the greatest generality we could achieve.

In everyday reasoning, we may be moderately concerned with sufficiency, but seldom with necessity. And even in the matter of sufficiency, we are usually very pragmatic. Gerhard Fischer has provided the following example, which comes from a Designer's-Assistant system that helps architects design kitchens. The system examines a design in which the architect has not placed the kitchen sink at a window, and critiques it, reporting to the architecture that the placement is undesirable. When the architect asks why, it replies: "The water pipes are usually installed with the expectation that the sink will be at the window, and are more accessible there. More important, the housewife spends a good deal of time at the sink, and finds it pleasanter to have a view from a window."

Now these are only two considerations out of many that must go into a kitchen design. In particular, they do not take into consideration the conflicting requirements that may be imposed on other fixtures and the problem of tradeoffs among them. Yet, they serve as "reasons" for placing the sink in a particular place, and it is precisely fragmentary reasons of this kind that we use in everyday life to justify most of our actions. We tend to accumulate reasons, no subset of which is wholly necessary, but we neglect many interactions, so that our reasons are almost never sufficient.

A principal reason why such an unrigorous procedure is satisfactory is that we take our actions in constant interaction with the environment. We take a step, notice new conditions and constraints that our reasoning ignored, and adjust our next step accordingly. In designing, we enter our tentative design decisions on a drawing, and the drawing reveals interactions that were ignored in our premises. If these interactions are important, we convert them into new constraints (new reasons), and modify the design to take account of them.

We human beings, except when using the formal tools of algebra (preferably with the help of a large computer) are not at all good in handling simultaneous equations. We solve them by successive approximations, finding an answer that seems to fit a small subset of the equations, then modifying it to fit others. One of the skills the professional designer acquires in any domain is knowledge of which constraints to satisfy first so that the plan can be carried to completion with a minimum of revision.

Qualitative Reasoning

Let these comments suffice, for the moment, to characterize reasoning in general so that we can now turn to *qualitative* reasoning. Roughly speaking, reasoning is qualitative when it leads to conclusions that do not specify "how much." Driving an auto employs a production that can be paraphrased: "If you press down on the accelerator, the car will go faster."

Leaving aside various picky details, such as why we speak of the accelerator as determining velocity and not acceleration, what can we do with such a vague inference? What we can do with it is to regulate the speed of our car to any desired accuracy, pressing on the pedal if it is going too slowly for us, easing up if it is going too fast. We need no numbers, just a calculus of more and less and equals, an ordinal calculus.

In the field of artificial intelligence, there has been a rash of proposals during the past decade for schemes capable of doing qualitative, ordinal reasoning about complex systems. A formal theory providing solid foundations for such schemes has, in fact, existed for nearly fifty years, having been developed by the economists Griffith Evans and Paul Samuelson for use in economic reasoning. In somewhat less explicit form, it was used even earlier in such domains of physics as mechanics and thermodynamics. In particular, it supported a form of analysis often called "comparative statics."

A great deal of everyday thinking takes the form of qualitative reasoning. It is the exception, rather than the rule, that we reason in terms of numbers or exact quantities of any kind. Even when we wish a numerical answer, we may first sketch out our analysis in qualitative terms, or check a quantitative result qualitatively to see whether it "makes sense."

I should like to show, primarily by means of examples, how qualitative reasoning works, to demonstrate its robustness when the variables are subjected to arbitrary monotonic transformations, to relate it to simple principles of the calculus and simple results in differential equations, and finally, to relate it to reasoning about diagrams and their corresponding mental pictures.

Comparative Statics

A great many reasoning problems have the following general form: A system is in equilibrium, or in some sort of steady state. The exact position of equilibrium depends upon some system parameters, which we may regard as exogenous variables whose values are determined by mechanisms independent of the system itself. The value of one or more system parameters is changed -- is increased, say -- and after the distrubance the system settles into a new equilibrium. We wish to infer how various system variables will change from the original equilibrium to the new equilibrium, whether they will increase or decrease or remain constant.

Often questions like this can be answered even if we do not know the numerical values of the variables in the system or the parameters, but only their signs, that is, in purely qualitative terms. Consider a simple example from classical price theory in economics. The quantity of a good that suppliers will offer varies with the price, larger quantities being offered at higher prices. The quantity of the same good that buyers will purchase also varies with the price, smaller quantities being purchased at higher prices. There is a price, the equilibrium price, at which the quantity supplied equals the quantity demanded.

Now suppose that a tax is levied on the good. The cost of producing it will be increased by a corresponding amount, and the price at which any specified quantity of the good will be offered by the suppliers will be increased by the same amount. (If we draw a graph with price as the vertical axis, the entire supply curve -- the schedule showing the prices at which different quantities will be offered -- will be displaced upward by the tax.) The amounts that buyers will purchase at any given price will remain unchanged.

Presented with these assumptions, it is not too difficult to conclude that the new equilibrium price (the price at which the quantity offered will equal the quantity demanded) will be higher than the previous price, and that the quantity exchanged at the new price will be smaller than the quantity in the previous equilibrium. One way to carry out the reasoning is to argue that suppliers will demand a higher price for the quantity that has previously been exchanged, but fewer purchases will be made at this higher price. The reduction in demand will, in turn, exert a downward pressure on the price, and therefore the net increase in price will be somewhat less than the amount of the tax.

There is a notable gap in this reasoning: will the process actually converge to a new equilibrium? Perhaps the reduction in demand produced by the higher price will be so great that the price will then actually be carried *lower* than the previous equilibrium. And this price change may, in turn, produce a new reduction in the quantity offered by sellers, and so on. The observable fact is that we often carry out such qualitative reasoning without concerning ourselves with the problem of possible non-convergence. To deal with that problem requires a mores sophisticated form of qualitative reasoning which I will discuss a little later.

However, the general process I am describing, of propagating initial impacts through the system of variables, is familiar enough. We have a set of variables connected by a system of mechanisms. We consider a change in one of the mechanisms (i.e., in a parameter or an exogenous variable), and propagate that change through the variables via the mechanisms that connect them. For example, an increase in rainfall increases the wheat crop, which increases the amount of wheat offered on the market, which reduces the equilibrium price at which wheat is sold. Hence we confidently conclude that good growing weather will reduce the price of wheat.

This particular example is simpler than the previous one because it omits any subsequent feedback from the market price to the quantity that will be supplied (the latter being supposed to be determined by the weather). Propagation of the effect is straightforward, and does not raise any question of convergence or stability. Notice that the reasoning also depends on a ceteris paribus assumption, that only one parameter or exogenous variable has changed, the others remaining constant.

Monotonic Transformations

Since the kind of reasoning we have been illustrating depends only on relations of greater or less and not upon cardinal quantities, it remains valid even if we stretch or bend the scales on which we measure our variables, so long as the distortions are monotonic. For price in our economic example, we can substitute any other variable that is a positive monotonic function of price; and for quantity an arbitrary positive monotonic function of quantity. Our reasoning will be unchanged so long as the ordering of corresponding values of the variables is unchanged: the greater value remains the greater, and the smaller the smaller.

There is no difficulty in constructing a formal mathematics of ordinal variables, or what amounts to the same thing, of relations that are invariant under monotonic transformations. What is more important is that human beings can make the inferences implied by this mathematics quite readily, at least in simple cases. "Simple cases" means situations where an effect can be propagated linearly through the system without interactions among different paths of propagation.

Even if the signs of all the connections are known, the possibility of drawing unequivocal inferences disappears quite rapidly in the presence of interactions. The reason is fairly obvious: along one path of propagation the effect of a change in parameter upon a system variable may be positive, along another path the effect on the same variable may be negative. The net effect, the sum of the positive and negative effects, is not invariant under *monotonic transformations* of the variables. To draw definite conclusions under these circumstances, more information about the variables, information on their cardinal values, is required. Not surprisingly, people do not reason very well in the presence of such interactions, which require them, in effect, to solve simultaneous equations.

Qualitative Treatment of Differential Equations

The problem of convergence that was raised above can sometimes been dealt with if we know something about the dynamics of the system we are dealing with. The equations of a system in equilibrium can be "dynamicized" by displacing the dependent variable of each to the right-hand side, and replacing it by its time derivative. That is, starting with $y = ax$, we first write $0 = ax - y$, and then $dy/dt = ax - y$. The equilibria of the system of differential equations constructed in this way is simply the original set of equations.

Of course, we should not write down these dynamic relations arbitrarily. Each one is an assertion about the mechanism that determines the rate of change of the corresponding dependent variable (the one whose derivative appears). The system consisting of $0 = ax - by$ and $0 = cx - dy$ can be dynamicized in two ways depending on which equation we choose to represent the mechanism regulating x and which to represent the mechanism regulating y. Which is the correct representation is not a matter of fiat, but a substantive question about how the system under consideration actually works.

Now if we can agree as to how to dynamicize a system of static equations, we gain by this process new information about the system, and new abilities to draw conclusions about its comparative statics. For we can solve the differential equations, obtaining their roots as functions of the system parameters. But these functions determine whether the roots correspond to a stable or an unstable equilibrium. If we *assume* stability, then we obtain new conditions on the parameters, which give us additional information about the net changes that will be produced.

Returning to our earlier example of the demand and supply relations, we now replace the supply function by a differential equation determining the rate of change in the quantity supplied as a function of the difference between the supply price for that quantity and the market price. The second differential equation determines the rate of change in the price as a function of the difference between the quantities supplied and demanded, respectively, at the current market price.

In the new dynamic system, the conditions for stability constrain the system parameters so as to guarantee convergence of price and quantity adjustments, and will determine the signs of the changes in equilibrium price and quantity produced by a change in a parameter or an exogenous variable. If we dynamicize the system in other ways, of course the stability conditions will be different.

It can be shown that the stability conditions are invariant under monotonic transformations of the variables. Hence, stability depends on relations of greater and less among variables, and not on the cardinal properties of the individual variables.

All of these modes of ordinal reasoning, in both the static and dynamic cases, can be given a graphical or diagramatic interpretation, and the reasoning can be carried out either by reference to an actual diagram on paper or one visualized in the "minds eye." Hence, when we see people reasoning ordinally, we cannot be sure whether it is being accomplished in some symbolic way or by visualization. We turn next to a brief discussion of the diagramatic representation.

Reasoning About Diagrams

In analytic geometry we all learned to convert equations into geometrical curves or surfaces, and geometrical figures into the equations of their boundaries. Thus, the example given earlier of the effect of a tax on equilibrium price and quantity of a commodity can be represented by two curves, a supply function and a demand function, on a plane with price and quantity as the axes. The equilibrium point is the intersection of the two curves, and the effect of the tax is represented by a parallel displacement of the supply curve, producing a new intersection. The conclusions we reached earlier by "propagation" can be reached simply by observing the properties of the diagram, that is, the relative positions of the equilibria before and after displacement.

To handle diagramatically the questions of convergence and stability requires a more complex diagram. When we dynamicize the static system, the differential equations determine possible paths of the system in the P-Q plane. Such a path is uniquely determined for each initial point, and the paths cannot cross (since the system cannot proceed along two different paths from the same state). The collectivity of such paths is called the *direction field* for the system, and from the direction field, we can infer how the system will move from any starting point -- specificially whether it will move to an equilibrium.

Consider against, our commodity market (Figure 1). The static equations for the demand curve and supply curve correspond to the states in which the price and the quantity, respectively, are in equilibrium, and therefore where dp/dt = 0 and dq/dt = 0, respectively. Hence, the paths will be horizontal where they cross the former curve and vertical where they cross the latter. It can be seen that if the supply curve is steeper than the demand curve (demand more elastic than supply), the path will converge to the (stable) equilibrium point at the intersection of the supply and demand curves. If the supply is more elastic, the paths will diverge from the (unstable) equilibrium.

A little further consideration shows that these properties are invariant under stretching or contraction of the two axes -- under monotonic transformation of the variables. Hence our diagrammatic analysis exactly parallels the algebraic analysis in terms of ordinal variables.

Conclusion

Now we have come full circle back to the topic of visualization and the role of the mind's eye and of external visual displays in human thinking. I have tried to survey some of the main tools and processes that seem to be implicated in everyday reasoning -- the kind that carries us through the day, dealing with problems as they arise.

The thinking I have described does not look at all like formal logic, and only a little like mathematics. It makes use of a great multitude of inference rules, which are not tautological rules of logic but incorporate much real-world knowledge. It appears to be remarkably unconcerned with questions of sufficiency and necessity. When it deals with quantities, as it often must, it usually handles primarily their ordinal rather than their cardinal properties. For most people, at least, it makes great use of diagramatic representations, or mental diagrams in the mind's eye, which provide it with powerful inference processes. To compensate for its severe limitations in handling simultaneous relations, it proceeds by successive approximations, and halts when it has satisficed.

By the standards of formal logic, it is a jerry-built structure. But it gets us through the day.